arXiv:1609.09170v1 [q-bio.PE] 29 Sep 2016

Inference of epidemiological parameters from household stratified data

James N. Walker[a], Joshua V. Ross[a], Andrew J. Black[a,*]

[a]*School of Mathematical Sciences, University of Adelaide, Adelaide, SA 5005, Australia*

Abstract

We consider a continuous-time Markov chain model of SIR disease dynamics with two levels of mixing. For this so-called stochastic households model, we provide a method for inferring the model parameters—governing within-household transmission, recovery, and between-household transmission—from data of the day upon which each individual became infectious and the household in which each infection occurred. This is achieved by calculating an approximation for the likelihood that is evaluated using a novel combination of numerical methods. This allows us to use Bayesian Markov chain Monte Carlo to calculate posterior distributions for all parameters and hence the household reproduction number and the early growth rate of the epidemic. Our methodology is motivated by the analysis of first few hundred study data.

Keywords: household model, Markov chain, Bayesian inference, between-household transmission.

1. Introduction

First few hundred (FF100) studies are data collection exercises carried out in the early stages of pandemic influenza outbreaks [1, 2, 3, 4]. The aim of these is to characterise a novel strain to determine its impact and hence inform public health planning [5, 6]. FF100 studies involve the collection of data from households where one person is confirmed to be infected. The members of the household are surveilled to identify their time(s) of symptom onset and the study is continued until the first few hundred cases have been observed, or adequate characterisation has been achieved. Households are the primary unit of observation because they are convenient to surveil—in contrast to more general contact tracing—and a large fraction of transmission occurs within the household [7].

Stochastic models, where the population are split into households with different rates of mixing within and between households, are a natural framework

*Email: andrew.black@adelaide.edu.au

to understand such data [8]. Recent work inferred *within-household* epidemic parameters from this type of household stratified data [9, 10]. In [9], inference is performed using a Bayesian MCMC framework, with exact evaluation of the likelihood, returning posterior distributions for all parameters of interest and hence the within-household reproductive ratio, R_0. In this paper we present a method for performing inference for a full Markovian SIR household branching process and hence, additionally, inference of the between household transmission parameter. This then allows us to estimate the household reproductive number, R_*, and the early growth rate of the epidemic, r, which are of importance to public health response [5, 6].

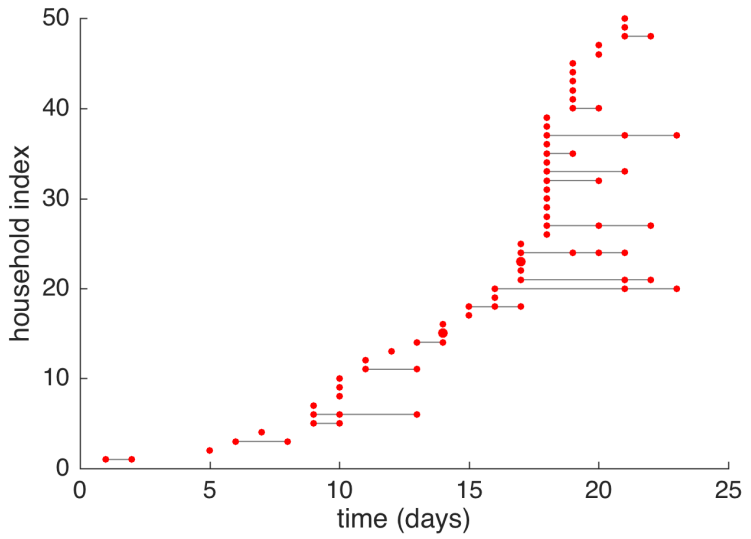

Figure 1: A realisation of the SIR household model (described in Section 2) with households of size 3. The times of symptom onset, binned into days, in the first 50 infected households at the beginning of an epidemic outbreak. The size of points corresponds to the number of infections on that day.

The data we assume to be available are illustrated in Figure 1; we observe only the times at a daily resolution when individuals become infected, which is assumed to coincide with symptom onset. The main challenge in this, as with many similar models, is that the likelihood is infeasible to compute due to the size of the model state space. Our approach is to carefully consider the dynamics and structure of the problem to allow us to derive an approximation to the exact likelihood that can be evaluated using a novel combination of numerical methods (matrix exponential methods [11], stochastic simulations [12] and numerical convolutions). This allows us to use the Metropolis-Hastings algorithm to compute posterior distributions for all the parameters of interest and hence R_*, the household reproductive number, and r, the early growth rate [13].

There are two main assumptions underpinning our work. The first is that we can approximate the early time behaviour of the epidemic as a branching process where only a single introduction to each household is possible. This is a very mild assumption and we would expect data collected in the early stages of an outbreak, say from an FF100 study, to conform to this reasonably closely. The second, more technical assumption we make, is that we can replace certain random variables that arise in the problem with their mean values. We argue that this provides a very good approximation and the final posteriors that we compute show excellent convergence to the true model parameters as the amount of household data is increased.

Two approaches that are often used for models with intractable likelihoods are particle marginal MCMC [14, 15] and approximate Bayesian computation (ABC) [16]. These methods both rely on simulation of the underlying model which is typically computationally cheap. Unfortunately, neither of these approaches are practical for this type of data because exact simulation of the process, although possible, is both complicated and expensive. The only other work we are aware of for performing inference on this type of branching process data is based on the expectation-maximisation (EM) algorithm (see [17] and references therein). This has the drawback of of being a frequentist, as opposed to a Bayesian, method. As far as we know, the method presented in this paper is the first for performing Bayesian inference on branching process data.

1.1. Overview of the paper

The paper is organised as follows. The Markovian household model and associated threshold quantities are introduced in Section 2. The data we assume to be available, and how we structure it for inference, are discussed in Section 3. We show how the likelihood can be decomposed into two parts related to the within- and between- household dynamics. In Section 4 we use an approximation to allow us to calculate the likelihood of seeing a given number of newly infected households over a day. This calculation depends on evaluating a number of conditional expectations relating to the dynamics of a single household. These calculations are detailed in Section 5 along with calculation of the part of the likelihood due to the within-household dynamics. Section 6 shows our results in the form of posterior distributions for the parameters of the model and demonstrates how these converge to the true parameters as more data is included from longer observations of the epidemic. We conclude the paper in Section 7 with a discussion of our method, its weaknesses and possible extensions. We also discuss the differences between our methodology and other potential inference methods.

2. Model and public health measures

The dynamics of the epidemic are modelled as a continuous-time Markov chain. Individuals are grouped into mutually exclusive households and make

effective contact at a high rate within households and at a low rate between households. As the data we wish to perform inference on is from the very earliest stages of an outbreak, we make the assumption that our model acts like a branching process at the household level, that is, we assume an infinite population size and hence that between-household infection only occurs into completely susceptible households. Households act independently after initial infection, hence we can consider the dynamics within each infected household, following their initial infection, in isolation from each other [8, 18].

In this paper, for simplicity, we will assume that all households are of the same size, N, and a basic SIR model for the within-household disease dynamics. Thus each individual is classified as *susceptible* to infection, S, *infectious* and able to infect susceptible individuals, I, or *recovered* and immune to the disease, R. As N is fixed, the state or configuration of a household can be specified by the number of susceptible and infectious individuals within the household (where $R = N - S - I$). Thus the state space for a single household is

$$\mathcal{S} = \{(s, i) \in \{0, 1, ..., N\}^2 : s + i \leq N\}.$$

The within-household dynamics are defined by the transitions that can occur and their corresponding rates. In a household with configuration (s, i), infectious individuals make effective contact within their household at rate β. The probability that effective contact within the household leads to an infection is $\frac{s}{N-1}$, hence the rate of within household infection is $\frac{\beta si}{N-1}$. Each infectious individual recovers at rate γ. The within household process can be defined in terms of its infinitesimal transition rate matrix, Q, given by

$$[Q]_{f(s,i),f(x,y)} = \begin{cases} \frac{\beta si}{N-1} & \text{for } (x, y) = (s - 1, i + 1),\ s \geq 1 \\ \gamma i & \text{for } (x, y) = (s, i - 1),\ i \geq 1 \\ -\frac{\beta si}{N-1} - \gamma i & \text{for } (x, y) = (s, i) \\ 0 & \text{otherwise,} \end{cases}$$

where $f : \mathcal{S} \to \{1, ..., |S|\}$ is a bijective map. The first infection within each household moves it into state $(N - 1, 1)$, at which point the within-household dynamics determine how the disease spreads within the household for the remainder of the epidemic.

We assume that between-household infection occurs due to homogeneous mixing of all the individuals in the population at rate α, thus the rate at which new households are infected is simply $\alpha I(t)$, where $I(t)$ is the total number of infected individuals in the population at time t. The model is initialised with a single infected household at a $U(0, 1)$ distributed time.

For this model we can identify the threshold parameter, R_*, which is a household (population level) reproduction number [8]. This is the expected number of households infected by a primary infectious household in an otherwise susceptible population of households; where a household is considered infectious

while it contains at least one infectious individual and a household is considered susceptible if it contains only susceptible individuals. It is one of at least five reproductive numbers that might be used when assessing the controllability of a disease in a community of households [19, 20, 21], but we adopt it herein as it is relatively easy to calculate and interpret. Let $\{X_t\}_{t\in\mathbb{R}^+}$ be the Markovian process that describes the state of an individual household from the time of its infection (i.e., the time of the first infection within the household). Let $I(k)$ be the function which returns the number of infectious individuals corresponding to state k. Then we have

$$R_* = E\left[\int_0^\infty \alpha I(X_t)\ dt\right],$$

where $X_0 = (N-1, 1)$ is the initial state of the process [8, 22, 23]. This can be calculated by solving a system of linear equations that depend on the parameters of the epidemic model [23, 24].

Also of interest from a public health perspective, is the early growth rate, r; this is also called the Malthusian parameter. Under the same conditions as above, this is defined as the unique solution to

$$E\left[\int_0^\infty \alpha I(X_t)e^{-rt}\ dt\right] = 1.$$

This can once again be evaluated efficiently [23].

3. Data and likelihood

Assume we observed an outbreak over some interval of time, $[0, T]$, and within that interval M households became infected. Each household is labelled by $i = 1, \ldots, M$ in the order that they became infected. Note that the process is only observed at a daily resolution (taken to be the end of each day) so the ordering within a day is arbitrary. The data available from each house are the cumulative number of infections at the end of each day since the first infection, where day t is defined as the interval $(t-1, t]$.

As each of the within-household outbreaks are assumed independent after their first infection, it is natural to specify this data in terms of two quantities: the days on which each household is infected and the time series of cumulative infection counts within each household, starting from their day of infection. More precisely, let ψ_t be the set of the labels (i) of the households that became infected on day t. Then let $\mathbf{w}^{(i)} = (w_j)^{(i)}$ be a vector where w_j is the cumulative number of infection events within the ith household, recorded at the end of day j, from the day of the households initial infection up to day T. Thus the data is completely specified by the sets $\{\psi_t\}_{t=1:T}$ and the vectors $\{\mathbf{w}^{(i)}\}_{i=1:M}$. These quantities are illustrated for a specific example in Figure 2. We also define $\Omega_t = \cup_{i=1}^{t-1}\psi_i$, which is the set of labels of households that became infected

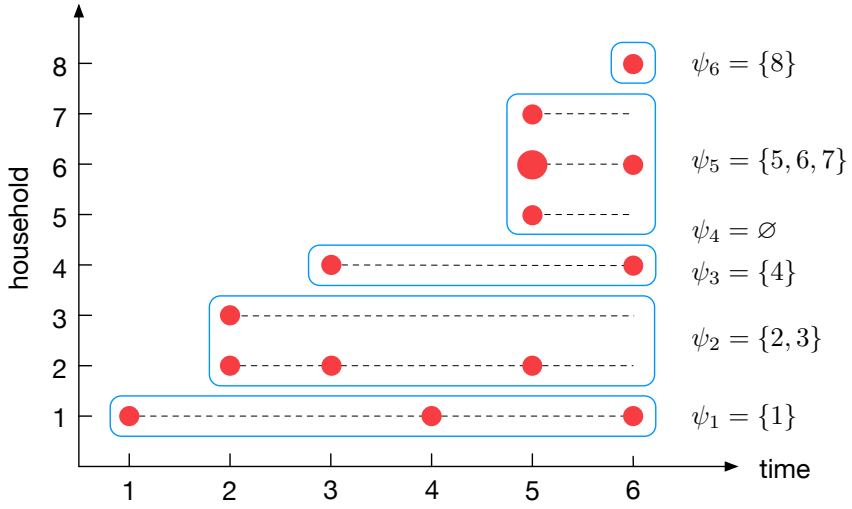

Figure 2: An illustration of how the data is structured for inference. An outbreak observed over $T = 6$ days, resulting in $M = 8$ households becoming infected. The red circles indicate the days on which new infections are observed and their size is proportional to the number of infections. The sets ψ_t indicate which households become infected on day t. Note that $\psi_4 = \varnothing$ indicates that no new houses were infected on day 4. The cumulative number of observed cases within each household, over the 6 days are: $\mathbf{w}^{(1)} = (1, 1, 1, 2, 2, 3)$, $\mathbf{w}^{(2)} = (1, 2, 2, 3, 3)$, $\mathbf{w}^{(3)} = (1, 1, 1, 1, 1)$, $\mathbf{w}^{(4)} = (1, 1, 1, 2)$, $\mathbf{w}^{(5)} = (1, 1)$, $\mathbf{w}^{(6)} = (2, 3)$, $\mathbf{w}^{(7)} = (1, 1)$ and $\mathbf{w}^{(8)} = (1)$.

before day t.

The initial condition for the process is taken to be a single household infected at a uniform time during the first day ($s \in (0, 1]$). The likelihood for $\theta := \{\alpha, \beta, \gamma\}$ can then be written as,

$$L(\theta) = \prod_{t=1}^{T} P\left(\{\mathbf{w}^{(i)}\}_{i\in\psi_t} \middle| \psi_t\right) P\left(\psi_t \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_t}\right); \tag{1}$$

here we have invoked the independence between $\left(\{\mathbf{w}^{(i)}\}_{i\in\psi_t} | \psi_t\right)$ and $\{\mathbf{w}^{(i)}\}_{i\in\Omega_t}$ due to the branching process assumption. Note that there is also a more subtle assumption of independence in this factorisation of the likelihood: even in the branching process approximation, the infection of new households does tell us information about the within-household processes—that there is at least 1 individual still infectious within the Ω_t households—but there is no easy way of incorporating this into the likelihood as is; in any case, the contribution to the likelihood of this will be small after more than two households have become infected.

As $\Omega_1 = \varnothing$, that is, there are no households infected before $t = 0$, the term

$$P\left(\psi_1 \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_1}\right) = P\left(\psi_1\right)$$

is determined by the initial condition. Further, households in ψ_t are identically distributed in the absence of within-household information. Thus their labels are arbitrary and only the number of households in ψ_t is relevant, that is

$$P\left(\psi_t \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_t}\right) = P\left(|\psi_t| \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_t}\right).$$

Thus we can factor the likelihood, Eq. (1), into two parts that are related to the within-household dynamics and between-household dynamics respectively. That is,

$$L(\theta) = L_w(\theta)L_b(\theta),$$

where

$$L_w(\theta) = \prod_{t=1}^{T} P\left(\{\mathbf{w}^{(i)}\}_{i\in\psi_t} \middle| \psi_t\right) \tag{2}$$

and

$$L_b(\theta) = \prod_{t=1}^{T} P\left(|\psi_t| \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_t}\right). \tag{3}$$

We refer to L_w as the within-household likelihood function and L_b as the between-household likelihood function. In Section 4 we detail how we calculate L_b and in Section 5, L_w.

4. Between-household likelihood, L_b

Each term in the product for the between-household likelihood, Eq. (3), is the probability that we observe $H_t := |\psi_t|$ new infected households on day t, given the data, over the time period $[0, T]$, for households that were infected before day t. We decompose H_t into two components, $H_t^{(1)}$ and $H_t^{(c)}$, such that $H_t = H_t^{(1)} + H_t^{(c)}$. The first component, $H_t^{(1)}$, is the number of the newly infected households on day t that are infected by a household in Ω_t, i.e., a household infected before day t. The second component, $H_t^{(c)}$, is the remaining number of newly infected households on day t, i.e., those that are infected by households that become infected on day t. We do not observe this demarcation in our data, but it assists us in the evaluation of the likelihood.

We start by considering the calculation of the probability mass function (pmf) of $H_t^{(1)}$, denoted $\mathbf{h}_t^{(1)}$, in Section 4.1. Then, in Section 4.2, we consider the evaluation of the pmf of $H_t^{(c)}$, $\mathbf{h}_t^{(c)}$. The required pmf of H_t, $\mathbf{h}_t$, is subsequently evaluated using efficient methods for calculating convolutions.

4.1. First generation of households, $H_t^{(1)}$

To calculate $\mathbf{h}_t^{(1)}$, the pmf of the number of first generation infected households, we note that on the first day of the epidemic there is only a single household infected at a $U[0,1]$ distributed time. Hence there is exactly 1 infected household in the first generation of households, so

$$P\left(H_1^{(1)} = 1\right) = 1.$$

For $t \geq 2$ we consider the rate at which the households in Ω_t infect new households. As we model the outbreak as a branching process, we assume that only completely susceptible households are infected, hence the instantaneous rate of infection at time $s \in (t-1, t]$ from the households in Ω_t is

$$\alpha \sum_{i \in \Omega_t} I\left(X_s^i\right),$$

where X_s^i is the state of household i at time s and $I(k)$ is a function returning the number of infectious individuals in a household in state k.

Thus the first generation of households are created as an inhomogeneous Poisson process, and conditioning on the information about the households in Ω_t, $\{\mathbf{w}^{(i)}\}_{i\in\Omega_t}$, we have

$$\left(H_t^{(1)} \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_t}\right) \sim \text{Poisson}\left(\alpha \sum_{i\in\Omega_t} \int_{t-1}^{t} I\left(X_s^i \middle| \mathbf{w}^{(i)}\right) ds\right). \tag{4}$$

Hence, we need to evaluate the distribution of

$$\Lambda_t := \alpha \sum_{i\in\Omega_t} \int_{t-1}^{t} I\left(X_s^i \middle| \mathbf{w}^{(i)}\right) ds.$$

However, this is expensive to compute, so instead we replace Λ_t in Eq. (4) with its expectation, which can be evaluated in a feasible manner. Precisely, we use

$$P\left(H_t^{(1)} = h \middle| \{\mathbf{w}^{(i)}\}_{i\in\Omega_t}\right) \approx \frac{e^{-E[\Lambda_t]} E[\Lambda_t]^h}{h!},$$

where

$$E[\Lambda_t] = \alpha \sum_{i\in\Omega_t} \int_{t-1}^{t} E\left[I\left(X_s^i \middle| \mathbf{w}^{(i)}\right)\right] ds. \tag{5}$$

As Λ_t is the force of infection over a short time period (a day) it should have relatively low variance, thus replacing Λ_t by its expectation may provide a reasonable approximation. In Section 5 we detail how the conditional expectations, Eq. (5), can be calculated using matrix exponential methods.

4.2. Subsequent generations of households, $H_t^{(c)}$

Recall that the number of newly infected households on day t is $H_t = H_t^{(1)} + H_t^{(c)}$. The first component, $H_t^{(1)}$, is the number of the newly infected households on day t that are infected by a household in Ω_t, i.e., a household infected before day t. The second component, $H_t^{(c)}$, is the remaining number of newly infected households on day t, i.e., those that are infected by households that become infected on day t.

We assume that the infection of the $H_t^{(1)}$ households are uniformly distributed over day t, and since their dynamics are independent, we have that $H_t^{(c)}$ is the convolution of $H_t^{(1)}$ random variables; we will use G to denote one of these random variables. Each of these random variables correspond to the size of a household branching process at time 1 that was initialised at a U$[0,1]$ time. The calculation of the pmf of G is once again computationally expensive, so here we choose to estimate this distribution using simulation [12]. As we are simulating over such a short period of time, and use the most efficient representation to minimise computational time, a large number of simulations, producing an accurate estimate, can be produced in a computationally-efficient manner.

Once we have estimated the pmf of G, we can calculate $\mathbf{h}_t$ from $\mathbf{h}_t^{(1)}$ as

$$\mathbf{h}_t = \mathcal{M}\mathbf{h}_t^{(1)},$$

where the convolution matrix, $\mathcal{M}$, is defined as follows. Let $\boldsymbol{\phi}$ be a column vector of the pmf of the random variable $G+1$. Then let

$$\mathbf{c}_i = \mathbf{c}_{i-1} * \boldsymbol{\phi}, \quad i \geq 2,$$

where '$*$' denotes a discrete convolution and $\mathbf{c}_1 = \boldsymbol{\phi}$. The matrix $\mathcal{M}$ is then given by

$$\mathcal{M} = [\mathbf{e}_1, \mathbf{c}_1, \mathbf{c}_2, \dots],$$

where $\mathbf{e}_1$ is a vector of 0s with the exception of a 1 in the first entry. The matrix M is truncated such that no probability needed for the calculation of the likelihood is lost. The calculation of this is not expensive, even for large matrices as the convolutions can be done using using discrete Fourier transforms [25]. In this paper we simply use the built in MATLAB function `conv()`, although other methods may provide computational gains, if required.

5. Single household dynamics, $E[\boldsymbol{\Lambda}_t]$ and $\boldsymbol{L}_w$

Recall, the evaluation of the pmf $\mathbf{h}_t^{(1)}$ for $t \geq 2$, corresponding to the number of first generation infected households on day t, requires the evaluation of the expected force of infection over day t from households infected prior to day t, $E[\Lambda_t]$. We begin by detailing the evaluation of $E[\Lambda_t]$, and then note how the

within-household likelihood, L_w, follows.

The computation of the $E[\Lambda_t]$ can be expressed in terms of integrals of the expected number of infectious individuals within each household in Ω_t, Eq. (5). As this expression is a sum over independent households we simplify our exposition, by detailing the calculation for a single arbitrary household in Ω_t, with observed data $\mathbf{w}$ (thus dropping the superscript 'i' notation for now). The independence also means we can rescale time within the household to begin at the start of the day of the first infection. Thus we need to calculate the expected number of infected individuals over each of the $|\mathbf{w}| = \omega$ days since the first infection within that house, i.e.

$$E\left[I\left(X_s\middle|\mathbf{w}\right)\right] = \sum_{k\in\mathcal{S}} I(k) P\left(X_s = k\middle|\mathbf{w}\right), \tag{6}$$

for all $s \in (t-1, t]$, where $t = 2, \ldots, \omega$. Note that as we are conditioning on the entire observed data within the household, $\mathbf{w}$, the random variable ω is implicitly conditioned on. That is, we are conditioning on knowing that the household became infected on day $T - \omega + 1$. In the remainder of this subsection all probabilities are conditioned on ω, but this is not written explicitly for concision.

The expectation Eq. (6) can be calculated efficiently, and hence the integral of the expectation to find the force of infection can also be calculated efficiently and accurately with Simpsons Rule, say. Our calculation is similar to that of the forward-backward algorithm [26], but is more involved as we need to calculate the expectations for all s, not just the discrete time points at which observations occur. First we define some quantities. As w_t is the total number of infections observed in the household by the end of day t (within-household time), X_t must be in a set of states such that $N - s = w_t$. These states are encoded by indicator vectors, $\mathbf{z}_t$, with 1s in entries corresponding to states where $N - s = w_t$ and zeros otherwise; these are either row or column vectors as required.

Define the row vector $\mathbf{f}_t$ as the 'forward' probabilities of the system, so the kth element is the probability the system is in state k at the end of day t, given the observed data up to t,

$$[\mathbf{f}_t]_k = P\left(X_t = k|\mathbf{w}_{(1:t)}\right).$$

These can be calculated in a recursive manner as follows:

$$\mathbf{f}_t = \frac{(\mathbf{f}_{t-1}e^{Q}) \circ \mathbf{z}_t}{(\mathbf{f}_{t-1}e^{Q}) \cdot \mathbf{z}_t}, \quad t = 2, \ldots, \omega,$$

where '$\circ$' is an element-wise vector product. The first vector, $\mathbf{f}_1$, is determined from the initial condition; let $\mathbf{v}$ be a probability vector with a 1 in the entry corresponding to state $(N-1, 1)$. As infection is introduced into each household at a $U[0, 1]$ distributed time on their day of infection, the distribution of X_1 is

given by

$$\mathbf{u} = \int_0^1 \mathbf{v} e^{Q(1-s)} ds.$$

Conditioning on w_1 gives $\mathbf{f}_1 = \mathbf{u} \circ \mathbf{z}_1 / \mathbf{u} \cdot \mathbf{z}_1$.

We also define the 'backward' probabilities, $\mathbf{b}_t$, with elements

$$[\mathbf{b}_t]_k = P\left(\mathbf{w}_{(t+1:\omega)} | X_t = k\right).$$

These are the probabilities of observing the remainder of the data given that the system is in state k at the end of day t. These can be calculated in a similar recursive way to the forward probabilities, but working backward from the final observation:

$$\mathbf{b}_{t-1} = e^Q (\mathbf{b}_t \circ \mathbf{z}_t), \quad t = \omega, \ldots, 2,$$

with $\mathbf{b}_\omega = \mathbf{1}$.

Applying Bayes' theorem to the pmf in Eq. (6) and using the Markov property we arrive at

$$P\left(X_s = k | \mathbf{w}\right) = \frac{P\left(X_s = k | \mathbf{w}_{(1:t-1)}\right) P\left(\mathbf{w}_{(t:\omega)} | X_s = k\right)}{P\left(\mathbf{w}_{(t:\omega)} | \mathbf{w}_{(1:t-1)}\right)}, \tag{7}$$

for $s \in (t-1, t]$. Using the law of total probability and the Markov property on the three probability expressions in Eq. (7) gives

$$P\left(X_s = k | \mathbf{w}_{(1:t-1)}\right) = \left[\mathbf{f}_{t-1} e^{Q(s-t+1)}\right]_k,$$

$$P\left(\mathbf{w}_{(t:\omega)} | X_s = k\right) = \left[e^{Q(t-s)} (\mathbf{b}_t \circ \mathbf{z}_t)\right]_k,$$

and

$$P\left(\mathbf{w}_{(t:\omega)} | \mathbf{w}_{(1:t-1)}\right) = \mathbf{f}_{t-1} \cdot \mathbf{b}_{t-1}.$$

Hence Eq. (6) can be expressed in a vectorised form as

$$E\left[I\left(X_s \middle| \mathbf{w}\right)\right] = \mathbf{i} \cdot \left(\frac{\mathbf{f}_{t-1} e^{Q(s-t+1)} \circ e^{Q(t-s)} (\mathbf{b}_t \circ \mathbf{z}_t)}{\mathbf{f}_{t-1} \cdot \mathbf{b}_{t-1}}\right), \quad t = 2, \ldots, \omega$$

where $\mathbf{i}$ is a vector whose elements are the number of infected individuals in each state.

This allows us to numerically evaluate $\mathbf{h}_t^{(1)}$; note that all matrix exponential calculations here can be expressed as $\left[e^Q\right]^a$, so we only need to compute the matrix exponential once per parameter set and take powers of the resulting matrix. Further, when numerically integrating, using a symmetric grid about $t - 1/2$ allows us to take advantage of the symmetry of $e^{Q(s-t+1)}$ and $e^{Q(t-s)}$, effectively halving the number of times we need to take powers of e^Q.

5.1. Within-household likelihood, L_w

Using the quantities calculated above, we can also calculate the within-household likelihood, L_w, described in Eq. (2). Let $\omega^{(i)}$ denote the length of $\mathbf{w}^{(i)}$. Note, under the branching process assumption, infected households act independently of each other, so their within household dynamics following their infection are independent. Hence,

$$\begin{aligned} L_w(\theta) &= \prod_{t=1}^{T} P\left(\left\{\mathbf{w}^{(i)}\right\}_{i\in\psi_t} \middle| \psi_t\right) \\ &= \prod_{t=1}^{T}\prod_{i\in\psi_t} P\left(\mathbf{w}^{(i)} \middle| i \in \psi_t\right) \\ &= \prod_{i=1}^{M} P\left(\mathbf{w}^{(i)} \middle| \omega^{(i)}\right). \end{aligned}$$

Let $\mathbf{f}_t^{(i)}$ and $\mathbf{z}_t^{(i)}$ denote the forward probability, and the state indicator vector on day t for household i, respectively. The probability of observing the data in each household is

$$\begin{aligned} P\left(\mathbf{w}^{(i)}|\omega^{(i)}\right) &= P\left(w_1\right)\prod_{t=2}^{\omega^{(i)}} P\left(w_t|\mathbf{w}_{(1:t-1)}\right) \\ &= \left(u\cdot\mathbf{z}_1^{(i)}\right)\prod_{t=2}^{\omega^{(i)}}\left(\mathbf{f}_{t-1}^{(i)}e^{Q}\right)\cdot\mathbf{z}_t^{(i)}, \end{aligned}$$

that is, the within-household likelihood is a product of the normalising constants for the forward probabilities. Thus the within-household likelihood is calculated as a by-product of the expectation calculations.

6. Results

Our inference methods are compared against 16 simulations with true parameter values $(\alpha, \beta, \gamma) = (0.32, 0.4, 1/3)$ and households of size $N = 3$. These parameters are chosen such that the average infectious period is three days (this is a typical infectious period for influenza), $R_0 = \beta/\gamma = 1.2$ and $R_* \approx 1.8$. The household size $N = 3$, noting the average household size in Australia is estimated to be 2.6 [27]. The number of realisations used to estimate the distribution of G was 10^3.

Inference is performed using a Bayesian Markov chain Monte Carlo (MCMC) framework, in particular we implement a Metropolis-Hastings algorithm in order to estimate the joint posterior distributions of our parameters [13]. The Metropolis-Hastings algorithm is run at various stages of the epidemic in order to show how the posterior distributions converge as more households become infected; the inference for each simulation is run after 50, 100, 200 and

300 households become infected. For each simulation, at each stage of the epidemic, 10^5 MCMC samples are obtained with a burn-in of 500 iterations. The Metropolis-Hastings algorithm is implemented with prior distribution for $\left(\alpha, \frac{\beta}{\gamma}, \frac{1}{\gamma}\right)$ of $U[0.05, 1] \times U[0.25, 4] \times U[0.25, 7]$ and the proposal distribution

$$X|Y \sim N\left(Y, \begin{bmatrix} 0.01 & 0 & 0 \\ 0 & 0.02 & 0 \\ 0 & 0 & 0.05 \end{bmatrix}\right).$$

Our results are displayed in terms of maximum a posteriori (MAP) estimates of the model parameters, (α, β, γ), and MAP estimates of key epidemiological parameters (R_*, r) and joint posterior density estimates of (R_*, r). All kernel densities were estimated using the freely available MATLAB packages `kde2.m` and `akde.m` [28, 29].

From Figure 3, we observe that MAP estimates begin negatively biased for all parameters and converge towards the true parameter values as more data is obtained. For all parameters the variability of the MAP estimates appear to decrease as more data is obtained. It should be noted that these box plots make the correlation structure of the parameters unclear; this is not presented here as the dimension of the parameter space makes the correlation structure difficult to display. In Figure 4 we observe that both r and R_* are over estimated initially, but converge towards the true values as more data is obtained. The posterior density estimate of (R_*, r) for a single simulation, displayed in Figure 5, shows that the posterior density initially has large variability, but the variability decreases as more households become infected.

7. Discussion

In this paper we have derived a method to calculate the likelihood for an SIR household branching process. This allows us to perform simple Bayesian MCMC inference to compute posteriors for both R_* and r, which are of importance for public health planning. These posteriors show excellent convergence to the true parameter values. In particular, they exhibit little variability after 200 households become infectious. It should be noted that due to computation time we have only presented results from 16 simulations and a single parameter set. This is the first study that we are aware of to estimate these quantities using early epidemic data and a Bayesian framework.

Another way to analyse the stochastic SIR household model is to cast it as a multi-type branching process [30, 31]. The theory of these is well developed and hence we can write down equations for many of the quantities we need in order to calculate a likelihood, but actually solving these is too inefficient for practical inference where the likelihood calculation is embedded in an MCMC scheme and hence needs to be repeated many times. Indeed the equations for

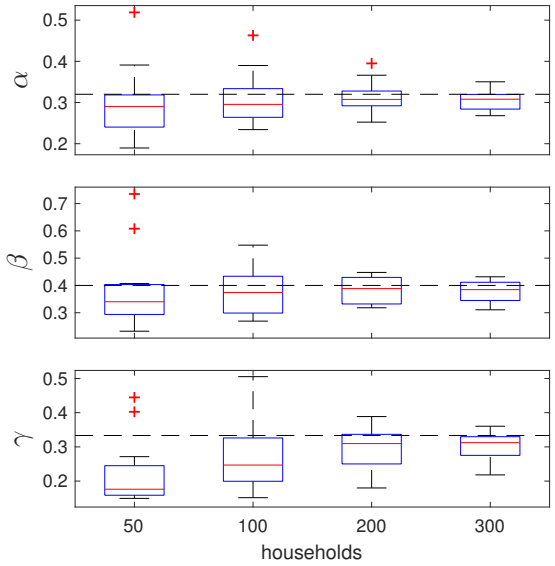

Figure 3: Boxplots of maximum a posteriori (MAP) estimates of (α, β, γ) from 16 simulations. MAP's are calculated from 3 dimensional kernel density estimates; each kernel density is estimated from 10^5 MCMC samples. The boxes from left to right are MAP's from inference based upon data with 50, 100, 200 and 300 infected households. Black dotted lines indicate the true parameter values at $(\alpha, \beta, \gamma) = (0.32, 0.4, 1/3)$.

the probability generating function for the full branching process are very simply stated, but their solution involves a multi-dimensional inversion [32]. As such, the approach we have taken is to factor the likelihood in a non-standard way, using a small number of well motivated assumptions. Our factorisation allows us to calculate its parts using a combination of numerical techniques; in particular, matrix exponential methods (Section 5) as well as stochastic simulations and numerical convolutions (Section 4.2). Each method is appropriate for the task and relatively efficient. For example the simulation to calculate the distribution of G can be programmed efficiently as there is no conditioning involved, so all the generated realisations can be used.

Still, this method is relatively expensive to run and there is room for improvement in many aspects of the procedure. In particular, no attempt has been made to parallelise any part of the algorithm. This would be relatively trivial as most of the calculations are independent of each other and hence this would provide a large speed-up. In a previous paper, [9], we also used a tree data structure to minimise the number of operations needed to calculate the within-household likelihood. A similar approach could be taken here to minimise the cost of the expectation calculations, as well as casting them as explicit path integrals that can typically be solved more efficiently [24].

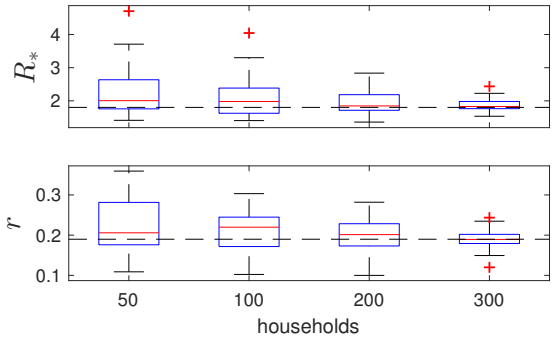

Figure 4: Boxplots of maximum a posteriori (MAP) estimates of (R_*, r) from 16 simulations. MAP's are calculated from 2 dimensional kernel density estimates; each kernel density is estimated from 10^5 MCMC samples. The boxes from left to right are MAP's from inference based upon data with 50, 100, 200 and 300 infected households. Black dotted lines indicate the true parameter values at $(R_*, r) \approx (1.803, 0.190)$.

There are a number of extensions that could be made quite easily to this methodology. An exposed period can be added to the model, but this changes the processes somewhat in that households are no longer observed at the time of first infection, but after the first individual becomes infectious (and displays symptoms). Thus we would need to track the distribution of exposed but not yet infectious households. One aspect that becomes easier with the addition of an exposed period is that longer chains of household infections become less likely on a given day. If the exposed period is sufficiently long with high enough probability (say, typically greater than 1 day) then we can approximate the distribution of newly exposed households with just a single generation. Another extension would be the incorporation of a realistic distribution of household sizes within the population. The expectation calculations would essentially remain the same, but the proportions of each size of household would need to be taken into account in the between-household likelihood calculation, in the simulations and the convolution procedure.

The biggest weakness of this work is that we assume perfect detection of infectious cases. Especially for diseases such as influenza, there can be a large fraction of asymptomatic cases and hence partial detection is the best that can be achieved. In previous work we have not made this perfect detection assumption, but instead assumed that there is some probability per case of detection [9]. Many aspects of this work could be extended to incorporate partial detection, but the largest challenge is modelling the distribution of currently unobserved households and how they contribute to the overall force of infection. The convolution approach may be appropriate here, especially given how fast this is using modern GPU hardware, but this is a topic for further research.

One approach that has the potential to analyse complex branching process

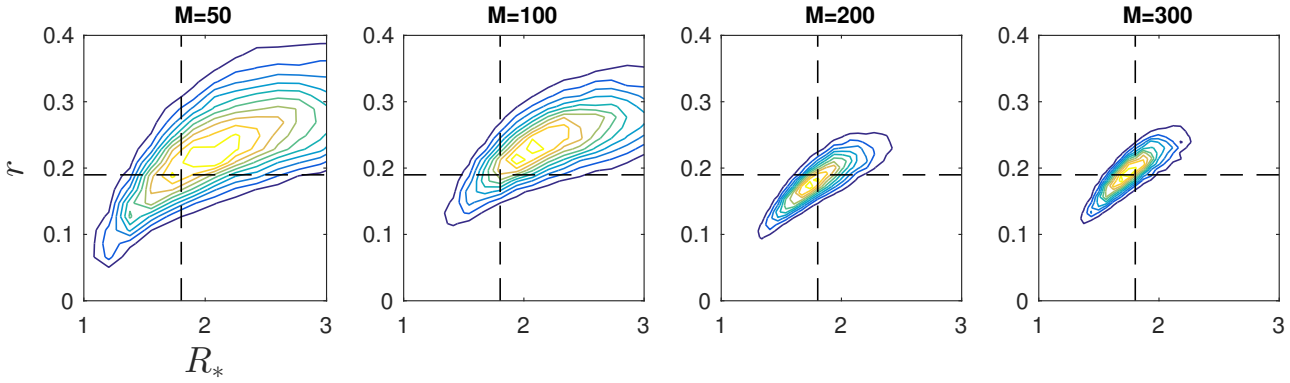

Figure 5: Contour plots of the joint posterior density of R_* and r from a single simulation. Each density is estimated from 10^5 MCMC samples. The panels from left to right are posteriors from inference based upon data with 50, 100, 200 and 300 infected households. The intersection of the black dotted lines indicate the true parameter values at $(R_*, r) \approx (1.803, 0.190)$.

data and partially observed data is data-augmented MCMC [33, 34, 35]. In this approach, all unobserved events are treated as unknowns to also be inferred within the MCMC routine; for the model considered in this paper these would be the infection and recovery times for each individual within each household. A data-augmented approach potentially allows much more flexibility in model choice and partial observation is also easy incorporated. The trade off of such an approach is that the MCMC scheme needed to sample from the joint distribution of parameters and unknown data is more complex and may need to run for longer to achieve proper mixing. Convergence can also be an issue when there is a large amount of missing data to be inferred [35]. Such an approach is clearly quite different to that adopted in this paper where we essentially consider all paths of the process at once for a given set of parameters. Our MCMC scheme is thus much simpler, but the calculations involved are more complicated, in particular the matrix exponentials. How these two approaches compare—and whether it is possible to combine aspects of them in one algorithm—is currently being investigated.

Acknowledgements

JNW acknowledges the support provided by the Dr Michael and Heather Dunne PhD Scholarship in Applied Mathematics. JVR acknowledges the support of the ARC (Future Fellowship FT130100254). AJB was supported by an ARC DECRA (DE160100690). All authors also acknowledge support from both the ARC Centre of Excellence for Mathematical and Statistical Frontiers (CoE ACEMS), and the Australian Government NHMRC Centre for Research Excellence in Policy Relevant Infectious diseases Simulation and Mathematical Modelling (CRE PRISM^2).

References

References

[1] Australian Department of Health and Ageing, Australian health management plan for pandemic influenza.

[2] Health Protection Agency England, "First Few Hundred" project, epidemiological protocols for comprehensive assessment of early swine influenza cases in the United Kingdom.

[3] E. McLean, et al., Pandemic (H1N1) 2009 influenza in the UK: clinical and epidemiological findings from the first few hundred (FF100) cases, Epidemiol. Infect. 138 (2010) 1531–41.

[4] A. B. van Gageldonk-Lafeber, van der Sande M. A., A. Meijer, I. H. Friesema, G. A. Donker, J. Reimerink, et al., Utility of the first few100 approach during the 2009 influenza A(H1N1) pandemic in the Netherlands, Antimicrob. Resist. Infect. Control 1 (2012) 30.

[5] J. M. McCaw, K. Glass, G. Mercer, J. McVernon, Pandemic controllability: a concept to guide a proportionate and flexible operational response to future influenza pandemics, Journal of Public Health 36 (2013) 5–12.

[6] C. Reed, M. Biggerstaff, L. Finelli, L. M. Koonin, D. Beauvais, A. Uzicanin, et al., Novel framework for assessing epidemiologic effect of influenza epidemics and pandemics, Emerg. Infect. Dis. 19 (2013) 85–91.

[7] A. Ghani, M. Baguelin, J. Griffin, A. J. van Hoek, S. Cauchemez, C. Donnelly, et al., The early transmission dynamics of H1N1pdm influenza in the united kingdom, PLoS Curr. 1 (2009) RRN1130.

[8] F. Ball, D. Mollison, G. Scalia-Tomba, Epidemics with two levels of mixing, Ann. App. Prob. 7 (1) (1997) 46–89.

[9] A. J. Black, N. Geard, J. M. McCaw, J. McVernon, J. V. Ross, Characterising pandemic severity and transmissibility from data collected during first few hundred studies (2016) Under review.

[10] F. Ball, L. Shaw, Estimating the within-household infection rate in emerging sir epidemics among a community of households, J. Math. Bio. 71 (2015) 1705.

[11] R. B. Sidje, Expokit: A software package for computing matrix exponentials., ACM Trans. Math. Softw. 24 (1998) 130–156.

[12] D. T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, J. Comput. Phys. 22 (1976) 403–434.

[13] W. R. Gilks, S. Richardson, D. J. Spiegelhalter, Markov Chain Monte Carlo in Practice, Chapman and Hall/CRC, 1995.

[14] C. Andrieu, A. Doucet, R. Holenstein, Particle Markov chain Monte Carlo methods (with discussion), J. R. Stat. Soc. B 72 (2010) 269–342.

[15] A. Golightly, D. J. Wilkinson, Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo, Interface Focus (2011) online.

[16] T. Toni, D. Welch, N. Strelkowa, A. Ipsen, M. Stumpf, Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems, J. R. Soc. Interface 6 (2009) 187–202.

[17] S. Hautphenne, M. Fackrell, An EM algorithm for the model fitting of markovian binary trees, Comput. Stat. Data Anal. 70 (2014) 19–34.

[18] A. J. Black, T. House, M. J. Keeling, J. V. Ross, Epidemiological consequences of household-based antiviral prophylaxis for pandemic influenza, J. R. Soc. Interface 10 (2013) 20121019.

[19] L. Pellis, F. Ball, P. Trapman, Reproduction numbers for epidemic models with households and other social structures I: Definition and calculation of R0., Math. Biosci. 235 (2009) 85–97.

[20] L. Pellis, F. Ball, P. Trapman, Reproduction numbers for epidemic models with households and other social structures II: comparisons and implications for vaccination. (2015) in prep.

[21] E. Goldstein, K. Paur, C. Fraser, E. Kenah, J. Wallinga, M. Lipsitch, Reproductive numbers, epidemic spread and control in a community of households., Math. Biosci. 221 (2009) 11–25.

[22] F. Ball, Stochastic and deterministic models for SIS epidemics among a population partitioned into households, Math. Biosci. 156 (1999) 41–67.

[23] J. V. Ross, T. House, M. J. Keeling, Calculation of disease dynamics in a population of households, PLoS ONE 5 (2010) e9666.

[24] P. Pollett, V. Stefanov, Path integrals for continuous-time Markov chains, J. Appl. Probab. 39 (2002) 901–904.

[25] R. G. Lyons, Understanding Digital Signal Processing, 3rd Edition, Prentice Hall, 2011.

[26] L. E. Baun, T. Petrie, G. Soules, N. Weiss, A maximisation technique occurring in the analysis of probabilistic functions of markov chains, Ann. Math. Statist. 41 (1970) 164–171.

[27] Australian Bureau of Statistics, Household and family projections, Australia, 2011 to 2036.

[28] Z. I. Botev, Kernel density estimation using MATLAB. Available at `http://au.mathworks.com/matlabcentral/profile/authors/878078-zdravko-botev`. Last checked 21 September 2016.

[29] Z. I. Botev, J. F. Grotowski, D. P. Kroese, Kernel density estimation via diffusion, Ann. Stat. 38 (2010) 2916–2957.

[30] K. Athreya, P. Ney, Branching Processes, Springer-Verlag, New York, 1972.

[31] K. S. Dorman, J. S. Sinsheimer, K. Lange, In the garden of branching processes, SIAM Review 46 (2004) 202–229.

[32] J. Abate, W. Whitt, Numerical inversion of probability generating functions, Operations Research Letters 12 (1992) 245–251.

[33] P. D. O'Neill, G. O. Roberts, Bayesian inference for partially observed stochastic epidemics, J. R. Stat. Soc. A 162 (1999) 121–130.

[34] N. Demiris, P. D. O'Neill, Bayesian inference for epidemics with two levels of mixing, Scand. J. Statist. 32 (2005) 265–280.

[35] T. J. McKinley, J. V. Ross, R. Deardon, A. R. Cook, Simulation-based Bayesian inference for epidemic models, Comput. Stat. Data Anal. 71 (2014) 434–447.